

time divided equally between course work and research. In the first year, the emphasis is on acquiring fundamental concepts and tools through course work and project involvement. During the second year, students implement and document a substantial AI application project.

III.A.3.6. Resource Operations and Usage

The following data give an overview of various aspects of SUMEX-AIM resource usage. There are 5 subsections containing data respectively for:

1. Overall resource loading data (page 53).
2. Relative system loading by community (page 54).
3. Individual project and community usage (page 57).
4. Network usage data (page 64).
5. System reliability data (page 64).

For the most part, the data used for these plots cover the entire span of the SUMEX-AIM project. This includes data from both the KI-TENEX system and the current DECsystem 2060. At the point where the SUMEX-AIM community switched over to the 2060 (February, 1983), you will notice severe changes in most of the graphs. This is due to many reasons briefly mentioned here:

1. Even though the TENEX operating system used on the KI-10 was a forerunner of the current Tops20 operating system, the Tops20 system is still different from TENEX in many ways. Tops20 uses a radically different job scheduling mechanism, different methods for computing monitor statistics, different I/O routines, etc. In general, it can not be assumed that statistics measured on the TENEX system correlate one to one with similar statistics under Tops20.
2. The KL-10 processor on the 2060 is a faster processor than the KI-10 processor used previously. Hence, a job running on the KL-10 will use less CPU time than the same job running on the KI-10. This aspect is further complicated by the fact that the SUMEX KI-10 system was a dual processor system.
3. The SUMEX-AIM Community was changing during the time of the transfer to the 2060. The usage of the GENET community on SUMEX had just been phased out. This part of the community accounted for much of the CPU time used by the AIM community. Since the purchase of the 2060 was partially funded by the Heuristic Programming Project (HPP), an additional number of HPP Core Research Projects started using the 2060, increasing the Stanford communities usage of the machine. And finally, the move to the 2060 occurred during a pivotal time in the community when more and more projects were either moving to their own local timesharing machines, or onto specialized Lisp workstations. It also was the time for the closure of many long time SUMEX-AIM projects, like DENDRAL and PUFF/VM.

Any conclusions reached by comparing the data before and after February, 1983 should be done with caution. The data is included in this years annual report mostly for casual comparison.

Also, it should be noted that monthly statistics are not available for this past year because of problems with the accounting program at this writing. The appropriate average data quantity for the year is shown instead for each month so the graphs appear to be "flat" in the area corresponding to the current period.

Overall Resource Loading Data

The following plot displays total CPU time delivered per month. This data includes usage of the KI-TENEX system and the current DECsystem 2060.




Figure 6: Total CPU Time Consumed by Month

Relative System Loading by Community

The SUMEX resource is divided, for administrative purposes, into three major communities: user projects based at the Stanford Medical School (*Stanford Projects*), user projects based outside of Stanford (*National AIM Projects*), and common system development efforts (*System Staff*). As defined in the resource management plan approved by the BRP at the start of the project, the available system CPU capacity and file space resources are divided between these communities as follows:

Stanford	40%
AIM	40%
Staff	20%

The "available" resources to be divided up in this way are those remaining after various monitor and community-wide functions are accounted for. These include such things as job scheduling, overhead, network service, file space for subsystems, documentation, etc.

The monthly usage of CPU resources and terminal connect time for each of these three communities relative to their respective aliquots is shown in the plots in Figure 7 and Figure 8. As mentioned on page 52, these plots include both KI-10 and 2060 usage data.

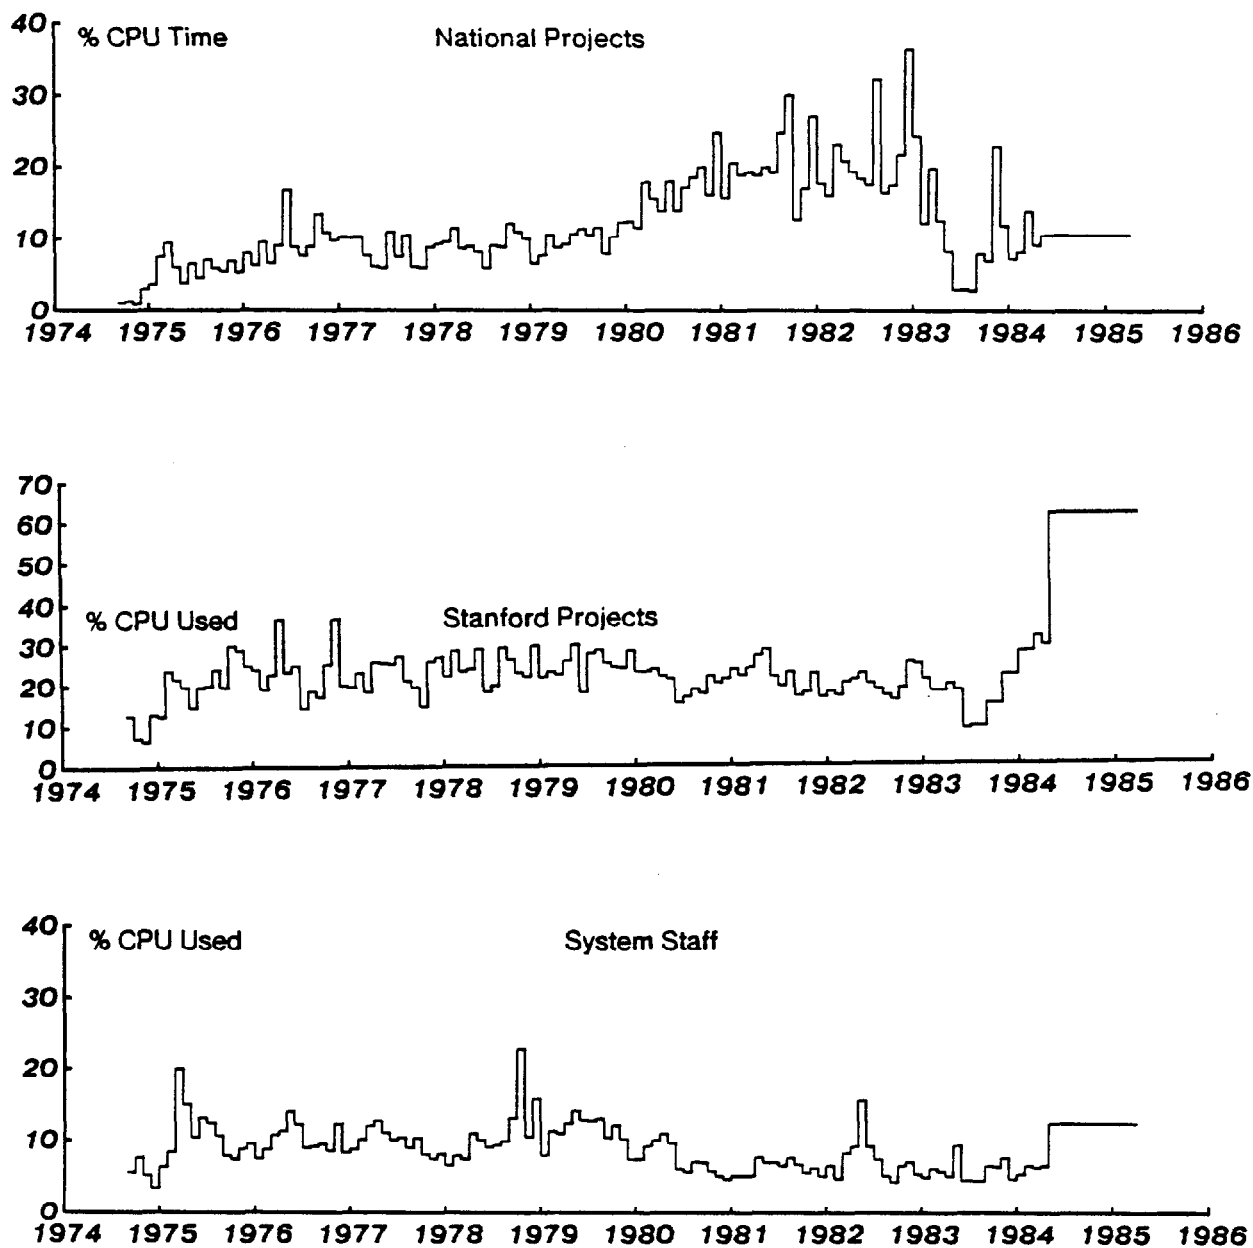


Figure 7: Monthly CPU Usage by Community

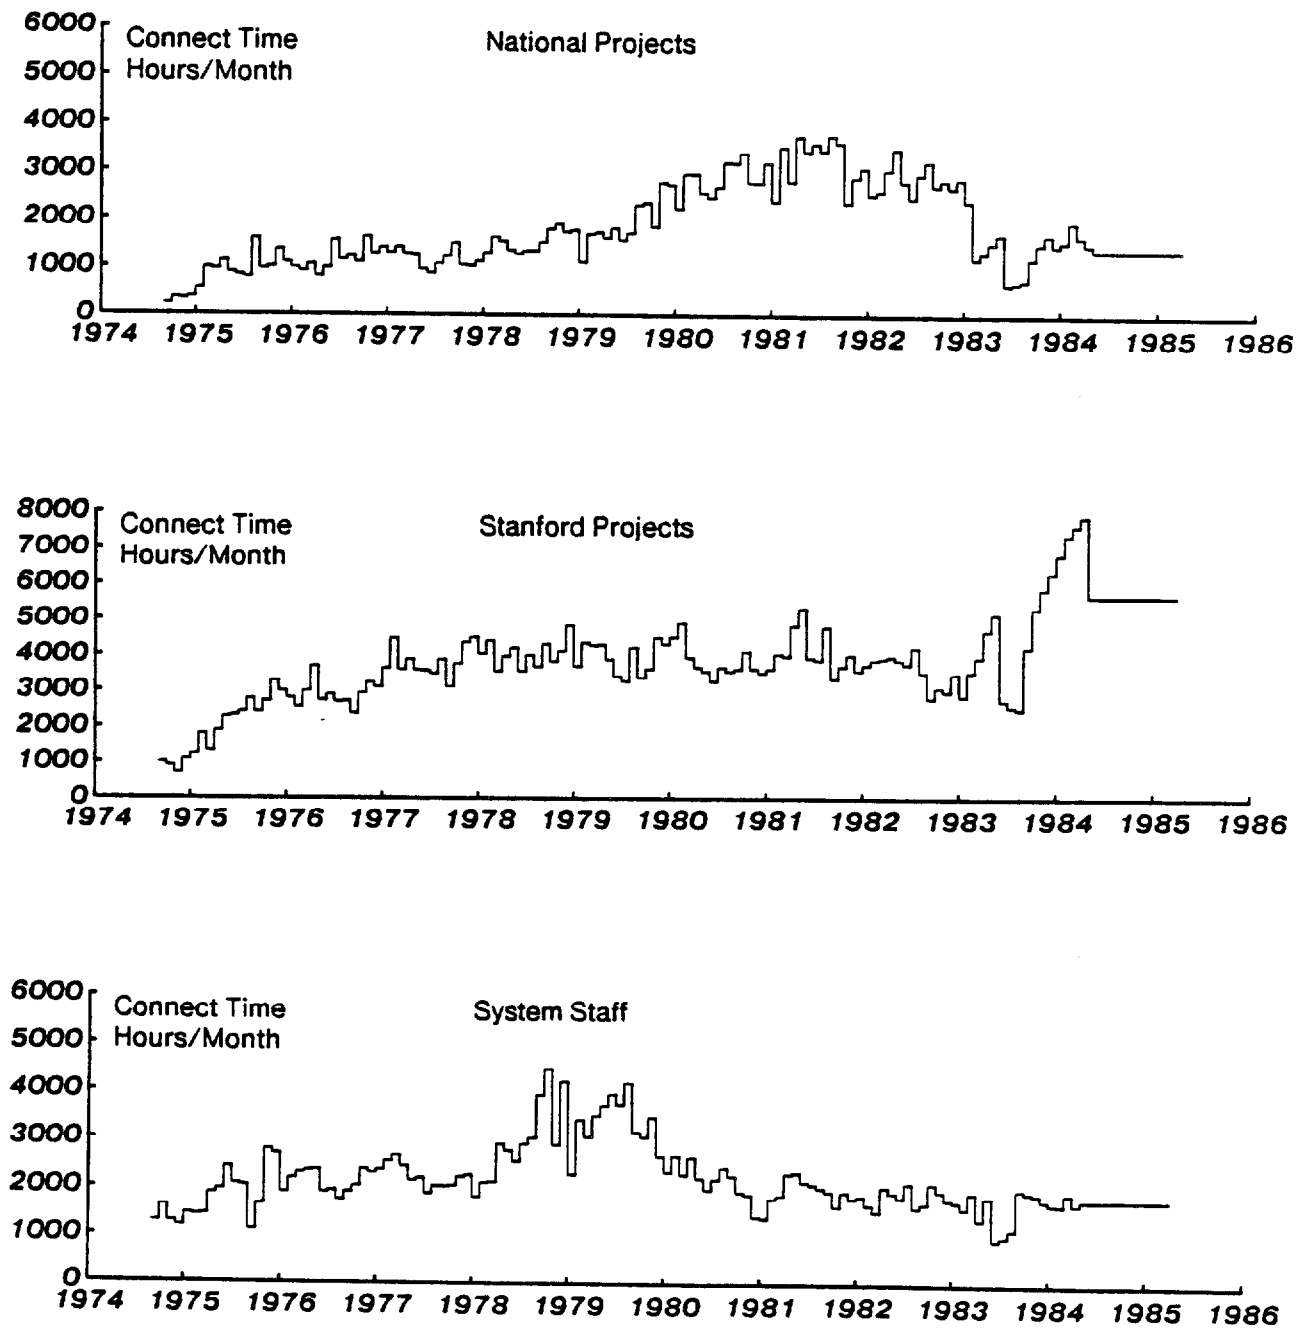


Figure 8: Monthly Terminal Connect Time by Community

Individual Project and Community Usage

The following histogram and table show cumulative resource usage by collaborative project and community during the past grant year. The histogram displays the project distribution of the total CPU time consumed between May 1, 1984 and April 30, 1985, on the SUMEX-AIM DECsystem2060 system.

In the table following, entries include a text summary of the funding sources (outside of SUMEX-supplied computing resources) for currently active projects, total CPU consumption by project (Hours), total terminal connect time by project (Hours), and average file space in use by project (Pages, 1 page = 512 computer words). These data were accumulated for each project for the months between May, 1984 and May, 1985.




Figure 9: Cumulative CPU Usage Histogram by Project and Community

Resource Use by Individual Project - 5/84 through 4/85

<i>National AIM Community</i>	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) CADUCEUS "Clinical Decision Systems Research Resource" Jack D. Myers, M.D. Harry E. Pople, Jr., Ph.D. University of Pittsburgh	86.72	1809.97	8028
2) CLIPR Project "Hierarchical Models of Human Cognition" Walter Kintsch, Ph.D. Peter G. Polson, Ph.D. University of Colorado	1.14	119.94	129
3) SECS Project "Simulation & Evaluation of Chemical Synthesis" W. Todd Wipke, Ph.D. U. California, Santa Cruz	45.14	5542.39	12230

4) SOLVER Project "Problem Solving Expertise" Paul E. Johnson, Ph.D. William B. Thompson, Ph.D. University of Minnesota	4.70	413.29	621
5) MENTOR Project "Medical Evaluation of Therapeutic Orders" Stuart M. Speedie, Ph.D. University of Maryland Terrence F. Blaschke, M.D. Stanford University	5.41	497.78	380
6) *** [Rutgers-AIM] *** Rutgers Research Resource Artificial Intelligence in Medicine Casimir Kulikowski, Ph.D. Sholom Weiss, Ph.D. Rutgers U., New Brunswick	0.62	57.29	196
7) AIM Pilot Projects	69.84	4292.54	3501
8) AIM Administration	0.42	57.86	673
9) AIM Users	27.88	3498.43	7135
	-----	-----	-----
Community Totals	241.87	16289.49	32893

<i>Stanford Community</i>	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) GUIDON-NEOMYCIN Project Bruce G. Buchanan, Ph.D. William J. Clancey, Ph.D. Dept. Computer Science	67.60	8225.93	6048
2) MOLGEN Project "Applications of Artificial Intelligence to Molecular Biology: Research in Theory Formation, Testing and Modification" Edward A. Feigenbaum, Ph.D. Peter Friedland, Ph.D. Charles Yanofsky, Ph.D. Depts. Computer Science/ Biology	238.64	8358.21	11392
3) ONCOCIN Project "Knowledge Engineering for Med. Consultation" Edward H. Shortliffe, M.D., Ph.D. Dept. Medicine	182.81	18869.06	16406

4) PROTEAN PROJECT Oleg Jardetzky School of Medicine Bruce Buchanan Computer Science Department	401.52	8539.01	13156
5) RADIX Project "Deriving Medical Knowledge from Time Oriented Clinical Databases" Robert L. Blum, M.D. Gio C.M. Wiederhold, Ph.D. Depts. Computer Science/ Medicine	33.23	2315.62	9168
6) Stanford Pilot Projects	277.71	6545.02	5092
7) Core AI Research	139.65	9447.97	10358
8) Stanford Associates	11.40	1030.22	1127
9) Medical Information Sciences	16.52	2561.42	974
	-----	-----	-----
Community Totals	1369.08	65892.46	70901

<i>KSL-AI Community</i>	CPU (Hours)	Connect (Hours)	File Space (Pages)
For funding details please see page 47			
1) Advanced Architectures	34.45	11070.95	3313
2) FOL	22.61	781.19	1522
2) Intelligent Agent	53.25	6934.73	3205
3) Pixie	12.98	1989.63	1072
4) KB VLSI	8.47	1275.64	927
5) KSL Management	114.18	21341.80	15597
6) DART	25.05	1497.89	12677
7) MRS	86.40	9298.69	1950
	-----	-----	-----
Community totals	357.39	54190.52	40263
<i>SUMEX Staff</i>	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) Staff	261.44	21450.55	17051
2) System Associates	26.84	1809.75	4744
	-----	-----	-----
Community Totals	288.28	23260.30	21795
<i>System Operations</i>	CPU (Hours)	Connect (Hours)	File Space (Pages)
1) Operations	775.69	69589.10	131640
	=====	=====	=====
Resource Totals	3032.31	229221.87	297492

(*) Award includes indirect costs.

System Reliability

System reliability for the DECsystem 2060 has significantly improved in this past period. We have had very few periods of particular hardware or software problems. The data below covers the period of May 1, 1984 to April 30, 1985. The actual downtime was rounded to the nearest hour.

Table 1 : System Downtime Hours per Month - May 1984 through April 1985

13	1	16	5	9	17	1	N/A	26	9	8	9
May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Jan	Feb	Mar	Apr

Table 2 : System Downtime Hours per Month - May 1984 through April 1985

Reporting period	:	364 days, 19 hours, 13 minutes, and 25 seconds
Total Up Time	:	359 days, 11 hours, 32 minutes, and 18 seconds
PM Downtime	:	1 days, 6 hours, 8 minutes, and 1 seconds
Actual Downtime	:	4 days, 1 hours, 33 minutes, and 6 seconds
Total Downtime	:	5 days, 7 hours, 41 minutes, and 7 seconds
Mtbf	:	3 days, 14 hours, 16 minutes, and 31 seconds
Uptime Percentage	:	98.89

Network Usage Statistics

The plots in Figure 10 and Figure 11 show the monthly network terminal connect time for the TYMNET and the INTERNET usage. The INTERNET is a broader term for what was previously referred to as Arpanet usage. Since many vendors now support the INTERNET protocols (IP/TCP) in addition to the Arpanet, which converted to IP/TCP in January of 1983, it is no longer possible to distinguish between Arpanet usage and Internet usage on our 2060 system.




Figure 10: TYMNET Terminal Connect Time




Figure 11: ARPANET Terminal Connect Time

III.A.4. Future Plans

Our plans for the next grant year (year 13) are based on the Council-approved plans for our 5-year renewal that began in August, 1980. The directions and background for much of this work were given in earlier progress report sections and are not repeated in detail here. Near- and long-term objectives and plans for individual collaborative projects are discussed in Section IV beginning on page 85.

Computing Resource Operation

The SUMEX-AIM resources -- mainframes, Lisp workstations, and networks -- provide crucial support for the AI research of our community. We will continue to operate these facilities for the most effective support of our users. We do not propose any substantial changes to the mainframe systems (DEC 2060, 2020, and shared VAX) but will continue to seek ways of minimizing maintenance costs and reliability.

We will continue to maintain operating system, language, and utility support software on our systems at the most current release levels, including up-to-date documentation. We also will be extending the facilities available to users where appropriate, drawing upon other community developments where possible. We rely heavily on the needs of the user community to direct system software development efforts.

Within the AIM community we expect to serve as a center for software-sharing between various distributed computing nodes. This will include contributing locally-developed programs, distributing those derived from elsewhere in the community, maintaining up-to-date information on subsystems available, and assisting in software maintenance.

Communication Networks

Networks have been centrally important to the research goals of SUMEX-AIM and will continue to be so for our increasingly distributed computing environment. Communication is crucial to maintain community scientific contacts, to facilitate shared system and software maintenance based on regional expertise, to allow necessary information flow and access at all levels, and to meet the technical requirements of shared equipment.

We have had reasonable success at meeting the geographical needs of the community through our ARPANET and TYMNET connections. These have allowed users from many locations within the United States and abroad to gain terminal access to the AIM resources and through ARPANET links to communicate much more voluminous file information. Since many of our users do not have ARPANET access privileges for technical or administrative reasons, a key problem impeding remote use has been the limited communications facilities (speed, file transfer, and terminal handling) offered currently by commercial networks. Commercial improvements are slow in coming and network delays have a major impact on remote projects -- mostly start-up pilot projects. We plan to continue experimenting with improved facilities as offered by commercial or government sources in the next year. We have budgeted for continued TYMNET service but will investigate alternatives as well, taking account of experiences with other national resources like BIONET.

Community Management

We plan to retain the current management structure that has worked so well in the past. We will continue to work closely with the management committees to recruit the additional high-quality projects which can be accommodated and to evolve resource allocation policies which appropriately reflect assigned priorities and project needs. We

expect the Executive and Advisory Committees to play a continuing role in advising on priorities for facility evolution and on-going community development planning in addition to their recruitment efforts. The composition of the Executive Committee will continue to represent major user groups and medical and computer science applications areas. The Advisory Group membership spans both medical and computer science research expertise. We expect to maintain this policy.

We will continue to make information available about the various projects both inside and outside of the community and, thereby, promote the kinds of exchanges exemplified earlier and made possible by network facilities.

The annual AIM workshops have served a valuable function in bringing community members and prospective users together. We will continue to support this effort. In July 1985, the AIM workshop will be hosted by the National Library of Medicine. We will continue to assist community participation and provide a computing base for workshop demonstrations and communications. We also will assist individual projects in organizing more specialized workshops as we have done for the DENDRAL and AGE projects.

We plan to continue indefinitely our present policy of non-monetary allocation control. We recognize, of course, that this increases our responsibility for the careful selection of projects with high scientific and community merit.

Training and Education Plans

We have an on-going commitment, within the constraints of our staff size, to provide effective user assistance, to maintain high-quality documentation of the evolving software support on the SUMEX-AIM system, and to provide software help facilities such as the HELP and Bulletin Board systems. These latter aids are an effective way to assist resource users in keeping informed about system and community developments and solving usage problems. We plan to take an active role in encouraging the development and dissemination of community knowledge resources such as the AI Handbook, up-to-date bibliographic sources, and developing knowledge bases. Since much of our community is geographically remote from our machine, these on-line aids are indispensable for self-help. We will continue to provide on-line personal assistance to users within the capacity of available staff through the MM and TALK facilities.

Core Resource Development

Our primary focus for core resource development will be in the area of Lisp workstations including improvements to the computing environment they offer, facilitating their interaction with each other, and enhancing their interaction with network services. This will include bringing up tools like electronic mail, text processing, file management, and others that we currently rely almost entirely on mainframe computers for. We will study problems of high-performance network protocol and file service for these workstations as well as general access to network printing facilities. We will continue the development of virtual network and graphics interfaces for the workstations so they can be more geographically accessible and so their total computing power can be exploited.

Core AI Research

Our basic AI research projects focus on understanding the roles of knowledge in symbolic problem solving systems -- its representation in software and hardware, its use for inference, and its acquisition. We are continuing to develop new tools for system builders and to improve old ones. In particular, we will focus on four areas with

immediate coupling to biomedical applications problems and on several others that may have future application. These include the Blackboard model of reasoning, constraint satisfaction systems, knowledge acquisition and learning, qualitative simulation, and other areas such as architectures for highly concurrent symbolic computation, a retrospective on the AGE blackboard tool, logic-based systems, self-aware systems, and the SOAR general problem-solving architecture.

III.B. Highlights

In this section we describe several research highlights from the past year's activities. These include notes on existing projects that have passed important milestones, new pilot projects that have shown progress in their initial stages, and some other special activities that reflect the impact and influence that the SUMEX-AIM resource has had in the scientific and educational communities.

III.B.1. Scholarly Publications

One of the important responsibilities of developers of a new technological area, such as artificial intelligence, is the scholarly assimilation and documentation of incremental progress. In addition to the numerous technical papers that have been published, 11 major books have been published from our community in the past 4 years:

- *Heuristic Reasoning about Uncertainty: An AI Approach*, Cohen, Pitman, 1985.
- *Readings in Medical Artificial Intelligence: The First Decade*, Clancey and Shortliffe, Addison-Wesley, 1984.
- *Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Programming Project*, Buchanan and Shortliffe, Addison-Wesley, 1984.
- *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World*, Feigenbaum and McCorduck, Addison-Wesley, 1983.
- *Building Expert Systems*, F. Hayes-Roth, Waterman, and Lenat, eds., Addison-Wesley, 1983.
- *System Aids in Constructing Consultation Programs: EMYCIN*, van Melle, UMI Research Press, 1982.
- *Knowledge-Based Systems in Artificial Intelligence: AM and TEIRESIAS*, Davis and Lenat, McGraw-Hill, 1982.
- *The Handbook of Artificial Intelligence*, Volume I, Barr and Feigenbaum, eds., 1981; Volume II, Barr and Feigenbaum, eds., 1982; Volume III, Cohen and Feigenbaum, eds., 1982; Kaufmann.
- *Applications of Artificial Intelligence for Organic Chemistry: The DENDRAL Project*, Lindsay, Buchanan, Feigenbaum, and Lederberg, McGraw-Hill, 1980.

III.B.2. The PROTEAN Project

The biomedical goals of the PROTEAN project, under Professors Jardetzky and Buchanan at Stanford, are to use techniques from artificial intelligence to help in the determination of the 3-dimensional structure of proteins in solution. Empirical data from nuclear magnetic resonance (NMR) and other sources may provide enough constraints on structural descriptions to allow protein chemists to bypass the laborious methods of crystallizing a protein and using X-ray diffraction to determine its structure. This problem exhibits considerable complexity. Yet there is reason to believe that AI programs can be written that reason much as experts do to resolve these difficulties [16].

In the last year, the PROTEAN project has moved from the "idea phase" to the "demonstration phase":

- A highly interdisciplinary research team has been assembled which epitomizes the spirit of the SUMEX community. They include faculty in medicine and computer science, research associates in computer science (one with an MD degree), one MSTP graduate student, and other graduate students in Bio-Engineering, Chemical Engineering, and Computer Science (one with a PhD in Chemistry).
- A problem-solving framework, named BBl, has been debugged and is running. Much of the code already existed, some from the AGE system, but during this year it was extended and put into a coherent package. It is designed to be general enough to work with constraint satisfaction problems of many kinds, but has only been tested to date on the protein structure problem. One of the important extensions is to make reasoning about control as explicit as reasoning about objects in the domain.
- A geometry system has been designed and a prototype version has been written. This system is "low-level" code that manipulates objects in a 3-dimensional coordinate system and answers questions about locations, overlap, orientations, etc. This system depends on a representation of relative locations of objects with respect to an anchor, and with respect to one another. For example, if HELIX-1 is posted as the anchor, then HELIX-2 may be placed relative to HELIX-1 and other objects may be placed relative to HELIX-2.
- A manually operated version of PROTEAN was developed to allow Prof. Jardetzky and members of his laboratory to step through their own procedures for using NMR data to solve protein structures. The program allowed them to refine their procedures, and also allowed us to understand the procedures well enough to define knowledge sources that would carry out the same operations without manual intervention.
- A qualitative solution was found for the LAC-Repressor Headpiece, a protein of 51 amino acids. This approximate structure describes the relative positions of the three alpha-helices relative to one another, but does not place random coils. The structure is not completely determined by the constraints inferred from the NMR spectrum, so we have developed a representation of allowed volumes for the helices relative to one another.
- The IRIS graphics terminal has been coupled with the reasoning program to allow us to display the partial structures defined at any stage in the reasoning. The link is currently from a Xerox D-machine through a VAX to the IRIS. Display code, in the language C, has been written that allows

display of allowed volumes (as "halos" surrounding objects) and manipulation of objects on the screen.

- Knowledge sources have been written for BB1 that control the reasoning about solving protein structures. These define the heuristics used by biochemists to decide, for example, on the secondary structure to use as an anchor (the largest one with the most constraints with other parts of the structure). Enough knowledge sources have been defined so far to allow PROTEAN to reason autonomously through the first three-quarters of the problem solving cycles that biochemists use for the qualitative structure of the LAC-Repressor protein.
- A program, named MARCK, has been written that aids in the definition of new knowledge sources. It "watches" what an expert does manually to find the points at which the expert's reasoning and PROTEAN's reasoning diverge. Then it uses the problem solving context to help construct a new knowledge source that would make PROTEAN's reasoning agree with the expert's in that context. MARCK has been successfully used to define many of the control knowledge sources now in PROTEAN.

III.B.3. Software Export

The SUMEX-AIM community has widely distributed both our system software and our AI tool software to other academic, government, and industrial research groups in the United States and abroad. This form of "publication" allows others to critique our results and build on those foundations. To date, our AI tool exports include:

GENET	Prior to the establishment of the BIONET resource at IntelliCorp, we distributed 21 copies of the DNA sequence analysis programs and databases for both DEC-10 and DEC-20 systems.
EMYCIN	A total of 56 sites have received the EMYCIN [4, 34] package for backward-chained, rule-based AI systems.
AGE	The AGE [25] blackboard framework system has been sent out to 35 sites in versions for several machines.
MRS	The MRS [9] logic-based system for meta-level representation and reasoning has been provided to 76 sites.
Other Programs	Smaller numbers of copies of programs such as the SACON [2] knowledge base for EMYCIN, the GLISP [27] system (now distributed by Gordon Novak at the University of Texas), and the new BB1 [14, 13] system have been distributed.

A number of other software packages have been licensed or otherwise made available for commercial development including DENDRAL (to Molecular Designs, Ltd.), MAINSAIL (to Xidak, Inc.), UNITS (to IntelliCorp, Inc.), and EMYCIN (to Teknowledge, Inc. and Texas Instruments, Inc).

In addition, our system programs such as the TOPS-20 file recognition enhancements, the Ethernet gateway and TIP programs, the SEAGATE AppleBus to Ethernet gateway, the PUP Leaf server, the SUMACC development system for Macintosh workstations, and our Lisp workstation programs are well-distributed throughout the ARPANET community and the respective user communities.

III.B.4. The MENTOR Project

The MENTOR (Medical EvaluationN of Therapeutic ORders) project is a transcontinental collaboration between Dr. Terry Blaschke at Stanford and Dr. Stuart Speedie at the University of Maryland. The MENTOR project was initiated in December 1983 as a pilot effort and has been funded by the National Center for Health Services Research since January 1, 1985. MENTOR's goal is to design and develop an expert system for monitoring drug therapy for hospitalized patients that will provide appropriate advice to physicians concerning the existence and management of adverse drug reactions. Today, information is provided to the physician in the form of raw data which are often difficult to interpret. The wealth of raw data may effectively hide important information about the patient from the physician. This is particularly true with respect to adverse reactions to drugs which can only be detected by simultaneous examinations of several different types of data including drug data, laboratory tests, and clinical signs.

In order to detect and appropriately manage adverse drug reactions, extensive medical knowledge and problem solving is required. An Expert System consultant on drug reactions could effectively gather the appropriate information from existing record-keeping systems and continually monitor for the occurrence of adverse drug reactions. Based on a knowledge base about drugs, it could analyze incoming data and inform physicians when adverse reactions are likely to occur or when they have occurred. The MENTOR project is an attempt to explore the problems associated with the development and implementation of such a system and to implement a prototype of a drug monitoring system in a hospital setting.

A number of independent studies have confirmed that the incidence of adverse reactions to drugs in hospitalized patients is significant and that they are for the most part preventable. Moreover, such statistics do not include instances of suboptimal drug therapy which may result in increased costs, extended length-of-stay, or ineffective therapy. Data in these areas are sparse, though medical care evaluations carried out as part of hospital quality assurance programs suggest that suboptimal therapy is common.

Other computer systems have been developed to influence physician decision making by monitoring patient data and providing feedback. However, most of these systems use relatively simple criteria for possible reactions and do not try to represent the complex medical decision making process involved. One might speculate that the lack of widespread acceptance of such systems may be due to the fact that their recommendations are often rejected by physicians.

The MENTOR system will use AI techniques to represent and reason about the complex of knowledge and data important to controlling adverse drug reactions in a monitoring and feedback system to influence physician decision-making. The initial effort has focused on the overall system design and work has begun on constructing a system for monitoring potassium in patients with drug therapy that can adversely affect potassium. Antibiotics, dosing in the presence of renal failure, and digoxin dosing have been identified as additional topics of interest.

III.B.5. Blackboard Model Research

Projects in the KSL, have experimented with many frameworks for building systems including rule-based (EMYCIN), frame-based (UNITS), logic-based (MRS), and blackboard-based (AGE) frameworks. We have also experimented with various methods of inference and control, including goal-directed, data-directed, and opportunistic reasoning. Of the paradigms we know about, the one that seems to offer the most flexibility is the blackboard model of reasoning.

It allows an arbitrary mixing of data-driven inference steps ("bottom up") with model-driven steps ("top down"). It allows a hierarchy of levels of abstraction in the on-going problem solution formation, from the most abstract (the global situation) to the least abstract (the supporting data or problem conditions). And it allows multiple sources of knowledge to provide the problem-solving links between these levels (i.e., information fusion).

Though the Blackboard framework was conceived at Carnegie-Mellon during the DARPA Speech Understanding project in the early 1970's, it has received much of its scientific and practical development by work in the Stanford Knowledge Systems Laboratory. The first development here was the HASP system for passive sonar signal understanding. Subsequent efforts involved experiments with scientific applications to x-ray crystallography, to planning, and in the development of the first software tool to assist knowledge engineers in constructing systems using the Blackboard framework (AGE).

The goal of our continuing research on blackboard systems is to improve the usability, the flexibility, and the inferential power of this framework for handling problems of hypothesis formation, signal understanding, constraint satisfaction and planning. This framework is also the organizing basis for our research on concurrent symbolic processing. We are implementing a new, domain-independent system called BBl [13, 14] that incorporates a full range of blackboard tools and we are making these notions concrete by building a substantial application system in the BBl framework in order to experiment with tradeoffs in the design. Specifically, BBl is the basis for the PROTEAN project which is attempting to build a program that infers tertiary structure of proteins from NMR data (plus knowledge of primary and secondary structure).

BBl, like earlier blackboard systems, is a domain-independent "blackboard control architecture" that solves problems through the actions of independent knowledge sources that record, modify, and link individual solution elements in a structured database (the blackboard) under the control of a scheduler. It expands upon the standard architecture in that:

- It provides an interpretable, modifiable representation for knowledge sources with more flexible means for triggering appropriate ones and support facilities for knowledge source creation, modification, and checking.
- Its blackboard representation permits dynamic assignment of attributes and values to objects on the blackboard and provides selective, demand-driven inheritance of attributes from linked objects, with local caching of results.
- It provides explicit reasoning about control -- the selection and sequencing of knowledge source actions -- with control knowledge sources that construct dynamic control plans out of problem-solving heuristics on a control blackboard. It provides a vocabulary and syntax for expressing control heuristics and a simple scheduler decides which domain and control knowledge sources to execute by adapting to whatever control heuristics currently are recorded on the control blackboard.